public static
if(securityPr
return secur
public static final String
the target system factory
"<service name=\""+TSF+"\"
"<interface class=\""
"<implementation class=\""
"\"/>"
"</service>"
ies;
the target system service
"<service name=\""+TSS+"\"
interface class=\""
implementation class=\""
"/>"
service>"
job management
service name=\""+JMS
interface class=\""
implementation
"\"/>"
"</service>"
e){
the storage
"<service name=\""
"<interface class=
"<implementation
"\"/>"

UNICORE

UNICORE

Table of Contents

1. Executive Summary
2. 10 reasons to use UNICORE
3. UNICORE Forum e.V.
4. Client Software/Web Portal
5. Service and Architecture
6. Data Management
7. Workflows

UNICORE

Executive Summary

Synopsis

UNICORE – A European Federation Software Suite – comes with a history of more than 20 years. Originally initiated in the Supercomputing domain, today UNICORE is a general-purpose federation software suite. In its recent version, UNICORE follows the latest standards from the Grid and Web Services world and offers a rich set of features to its users.

UNICORE is used in e-infrastructures of any nature and without limitations on the type of computing resource. Single PCs, cluster-systems and leadership HPC systems like the ones forming the PRACE Research Infrastructure and the Extreme Science and Engineering Discovery Environment (XSEDE) are supported as well as various storage types. Data management functionality includes high-performance file transfer and metadata handling. UNICORE's flexible security architecture enables many usage scenarios.

All UNICORE software is available as Open Source under BSD license from the UNICORE Website (**www.unicore.eu**), while the software repository is hosted on SourceForge.

Supported by the UNICORE Forum (**www.unicore.eu/unicore-forum**), the Jülich Supercomputing Centre (JSC) at Forschungszentrum Jülich co-ordinates the UNICORE Open Source activities and provides support for academic users of UNICORE.

Contact

For downloading the UNICORE software, taking a look at the source code, reading documentation and tutorials, or finding out about developments in the community, please visit the UNICORE web site **www.unicore.eu**

UNICORE

10 reasons to use UNICORE

Manager

1. Mature, well-proven federation software suite; widely used in European and national initiatives and infrastructures
2. Integrated solution for federations of HPC systems, Clusters and storage systems
3. Fast and competent support directly from the developers

User & Resource Provider

4. Easy to install and configure clients and services
5. Intuitive portal and graphical user interface as well as powerful command line client for seamless access to resources
6. Application integration mechanisms on the client, services, and resource level
7. Mature workflow support
8. High throughput data transfer

Developer

9. Source code (all Java) is available and easy to modify (BSD licensed)
10. Easily extensible with own developments with the option of integration in official releases

UNICORE FORUM

UNICORE Forum e.V.

Overview

The UNICORE Forum e.V. was founded in December 1999 by developers, leading European HPC centres, and supporting hardware vendors as a non-profit association to promote the development and distribution of the UNICORE Grid technology.

Membership is open to users and developers of Grid software as well as to hardware vendors. The UNICORE Forum e.V. has 15 members.

Details can be found at **www.unicore.eu/unicore-forum/members**.

The statutes can be found at **www.unicore.eu/unicore-forum/statute**.

Objectives

- foster the distribution and use of UNICORE
- organize workshops and booths at major conferences
- support presentations at conferences
- publish and maintain the specifications
- coordinate further development
- certify implementations and extensions

Membership

According to §9 of the statute membership is open to all organisations (private or public) which may contribute to the goals of the UNICORE Forum.

More details can be found at **www.unicore.eu/unicore-forum/membership**.

Technical Advisory Board

Objectives

- develop the future strategy and roadmap of UNICORE
- drive and monitor the open source development of UNICORE
- prepare technical proposals regarding UNICORE
- evaluate technical proposals regarding UNICORE
- advise the Board of the UNICORE Forum on request

Members

Krzystof Benedyczak (ICM – Warsaw University)
Richard Grunzke (Technische Universität Dresden)
Ivan Kondov (Karlsruhe Institute of Technology)
Gert Ohme (T-Systems SfR)
Jedrzej Rybicki (Forschungszentrum Jülich)
Bernd Schuller (Forschungszentrum Jülich)

UNICORE @ ISC 2017

Contact

UNICORE Forum e.V.
Chairman of the Board
c/o Karlsruhe Institute of Technology (KIT)
Steinbuch Centre for Computing
Prof. Dr. Achim Streit
Hermann-von-Helmholtz-Platz 1
76344 Eggenstein-Leopoldshafen
Germany

Email: unicore-forum-board@unicore.eu

UNICORE

Client Software

UNICORE comes with a set of clients covering the range from easy-to-use Web Portal via feature rich graphical and command line client to RESTful API.

The UNICORE Web Portal

The UNICORE Portal is the latest client in the portfolio. In contrast to other UNICORE clients, the portal is accessed by a standard web browser, no need to install and manage any special software.

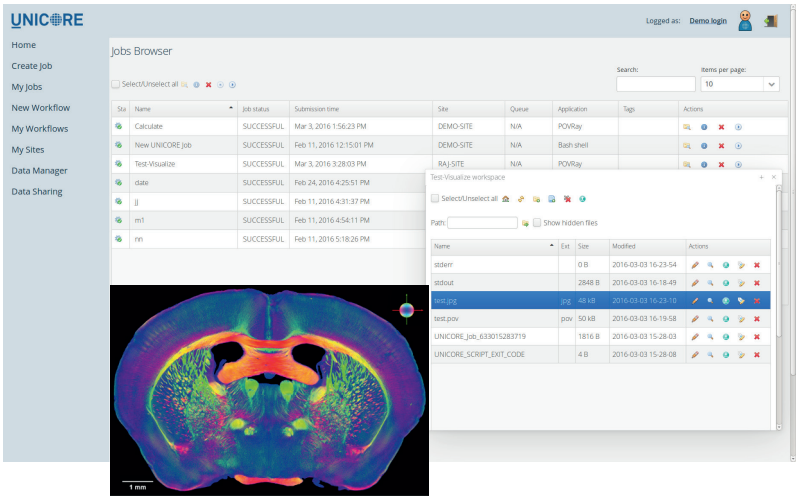

Authentication

The portal offers simple user management, registration of new users and integration with the Unity identity management solution. The traditional UNICORE authentication using X509 certificates and trust delegation is supported as well.

Features

It provides the following functionalities: user-friendly grid and sites browsing; easy job creation, submission, and monitoring; built-in workflow editor for basic use cases; fetching of job and workflow outcomes, as well as visualization of results; powerful data management, including remote file editing on the fly, access of remote storages, upload and download of local and remote files, and high–volume, high–throughput data transfer between local file system and remote storages. The newest features offer the upload and edit of workflow templates for easy parameter fitting as well as an integrated data sharing component.

The UNICORE Rich Client (URC)

The URC is an extensible Grid client framework based on Eclipse. The user interface allows for accessing and browsing any UNICORE based Grid. There are graphical editors for computational jobs and workflows that can be submitted to the appropriate execution services. Application specific graphical user interfaces can be easily integrated in order to obtain a high level view on the input and output data of scientific applications. Monitoring job and workflow execution is supported, as well as transferring remote files.

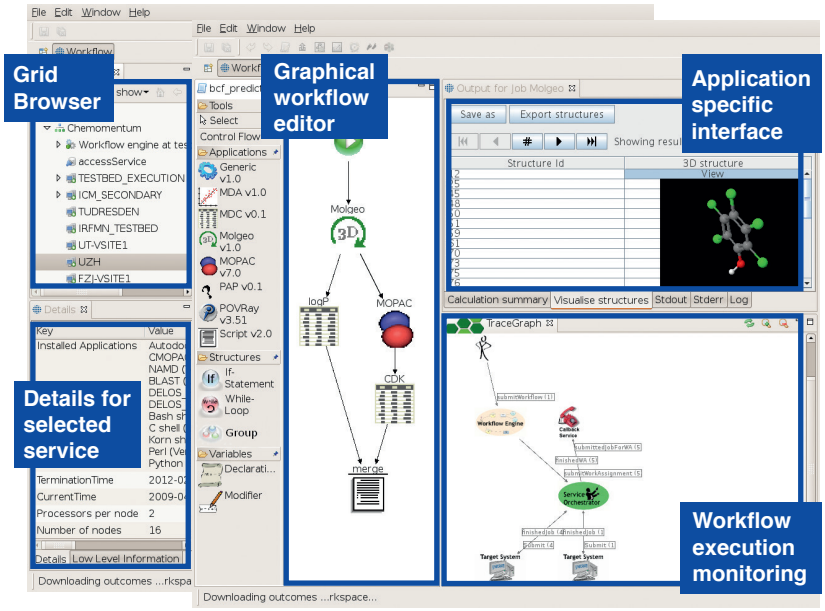

Grid browsing, workflow editing and execution monitoring with the UNICORE Rich Client

The workflow editor can be extended with new types of workflow structures and activities.

The UNICORE Commandline Client (UCC)

The UCC is a versatile commandline tool that enables users to access all features of the UNICORE service layer from a shell or scripting environment. Like the UNICORE Rich Client, it allows to run jobs and workflows, monitor their status and retrieve generated output. Additional features include a powerful batch mode for submitting multiple jobs in an automated fashion, listing and searching remote storages, transferring files, and administative tools for checking site.

The RESTful API

If you need a specific client for your environment you can exploit the RESTful programming interface to communicate with the UNICORE services.

UNICORE

Services and Architecture

Overview

UNICORE is a federation middleware suite for scientific and commercial domains providing seamless and secure access to computational and data resources.

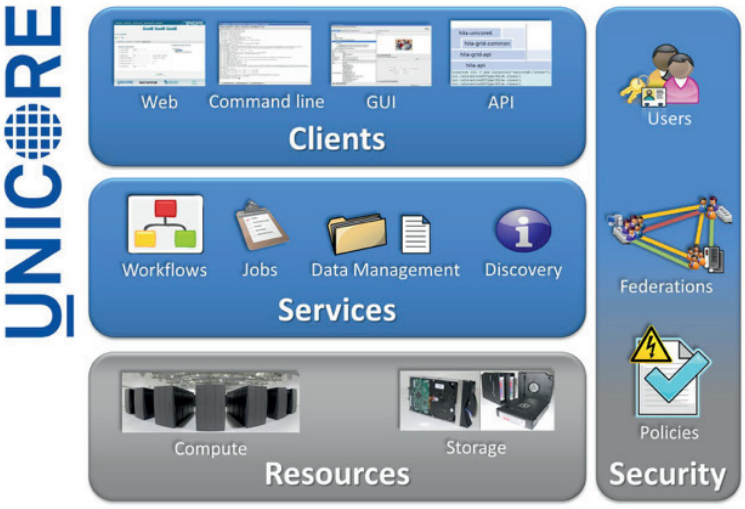

Portable

UNICORE is implemented in Java and Python, and runs on any recent operating system. A variety of high-performance computing and data storage solutions are supported out of the box.

Open

UNICORE uses SOAP and RESTful Web Services over HTTPS, as well as several other common open standards (SAML, XACML, JSDL, ...).

Extensible

UNICORE's modular architecture and open-source character provide for ease of extensibility.

Scalable

Multiple UNICORE installations can be combined to form a distributed, multi-organizational Grid allowing for thousands of jobs.

Features

The UNICORE architecture consists of the following major building blocks.

Services

UNICORE's service layer is implemented using the UNICORE Service Environment (USE). Out of the box, UNICORE provides a full set of Grid services (job submission and management, storage access, client-server and server-server data movement, metadata management), services for workflow support and service orchestration, as well as a service registry. Services are implemented using the Web Service Resource Framework (WSRF) standard, which essentially specifies stateful Web Services which are ideally suited for representing distributed computing resources. In parallel to the WSRF interfaces, RESTful interfaces are available.

Execution Management

The XNJS (Extensible Network Job Supervisor) is the processing engine at the heart of a UNICORE server. Used for managing job execution, it handles the mapping of abstract job definitions to the actual executables on the target systems.

Target System Interface

The Target System Interface (TSI) allows for invoking executables and accessing and modifying files on target systems. A typical UNICORE installation will use one of a set of different TSI implementations accessing the installed cluster resource management software such as Torque, Slurm and LoadLeveler. If required, the TSI can be adapted to the local environment in a straightforward manner.

Storages

Besides file systems accessed via the TSI, UNICORE can access various additional types of storage backends: Apache HDFS, Amazon S3 and CDMI. New storage types can be added easily.

Security

In order to establish a secure communications channel, UNICORE servers use X.509 certificates. UNICORE integrates with user management solutions such as Unity or LDAP, providing many options for authenticating users. User authorization uses fine grained, role-based access expressed via security policies.

Further Reading

www.unicore.eu/documentation/architecture

UNICORE

Data management

Overview

When realising complex scientific applications or industrial workflows, data management is one of the crucial ingredients. Consequently, UNICORE integrates powerful data and metadata management features.

Storage backends

The UNICORE SMS (storage management service) provides an abstract interface to a variety of backends. Currently, the following backends are supported:

- **File systems** can be accessed via the UNICORE Target System Interface (TSI). This is the default storage backend.
- **S3:** cloud storage accessible via the S3 interface (e.g. Amazon, OpenStack, ...). UNICORE can flexibly manage the required access keys.
- **Apache Hadoop HDFS:** the popular HDFS file system can be accessed via UNICORE as well, making it possible to combine Hadoop and HPC into application workflows.

Other backends can be added, as the implementations are pluggable.

Data transfer

The ability to move data efficiently from client to server, or between servers, is realised using file transfer services.

- **Multiple protocols:** depending on the storage back end, different file transfer protocols are available, and can be used transparently.
- **Server-server transfers:** when copying data from one server to another, the start time of the transfer can be controlled. Server-server transfers are managed by the server, so the client can go offline.
- **High performance data transfer:** the UNICORE file transfer protocol (UFTP) combines high performance with firewall-friendly configuration. Data connections are dynamically allocated and managed. Additional features include multiple TCP streams per transfer, optional data encryption and data compression, allowing to fine-tune UFTP to your requirements.

Metadata management

UNICORE has a built-in metadata system with a number of distinct features.

- **Metadata is stored** in hidden files right next to the data, which allows to keep data and metadata consistent and making backup and archiving easy.
- **Automated extraction:** UNICORE can automatically retrieve and index metadata from your files using the Apache Tika extraction framework.
- **Search engine:** indexing and querying the metadata using Apache Lucene, a powerful open-source Java search engine.
- **Search** only one given storage, or search multiple storages across the UNICORE federation.

Data-driven processing

UNICORE can be configured to automatically process files according to a set of user-defined rules. It's possible to run short scripts, batch jobs or to extract metadata.

Further Reading

www.unicore.eu/unicore/architecture

UNICORE

Workflows

Overview

Integrated workflow support is a major strength of UNICORE. In order to accomodate the wide range of requirements in scientific workflows, the UNICORE workflow system has been designed from the ground up for flexibility, scalability, extensibility and ease of use.

Workflows can be submitted and managed from all UNICORE clients. The graphical UNICORE Rich Client (URC) provides an extensive graphical workflow designer and advanced monitoring. Simple use cases can also be handled using the UNICORE Portal and the commandline client UCC.

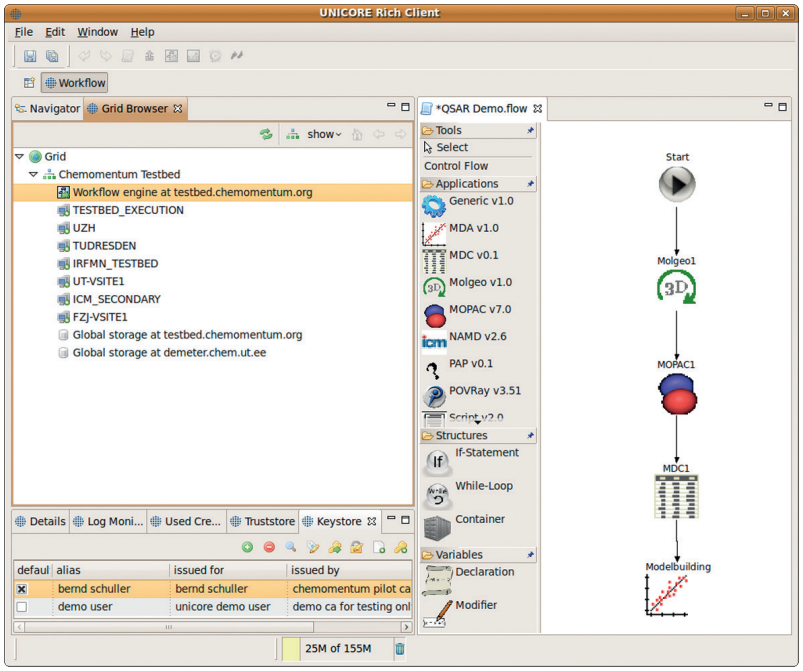

Workflow editing in the UNICORE Rich Client

Workflow features

The workflow is composed of individual application invocations and control structures. It supports common features such as parallel and sequential execution of tasks, loops, conditional transitions (e.g. if-else constructs) and workflow variables. The workflow system is ideally suited for scientific applications and parameter studies.

Architecture of the Workflow System

In the following the major building blocks of the workflow system are described.

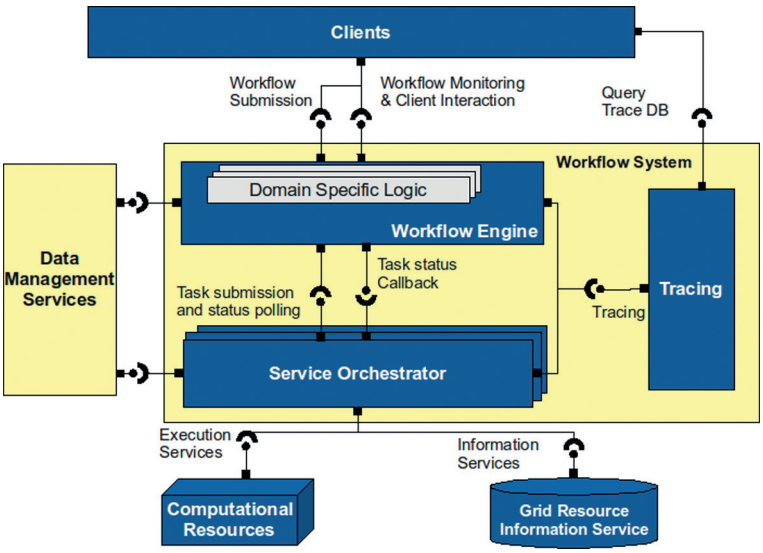

Component view of the UNICORE Workflow system

Workflow Engine

The Workflow Engine offers frontends that clients use to submit and manage workflows. It is responsible for processing the workflow logic and submitting tasks for execution to the Service Orchestrator.

Service Orchestrator

The Service Orchestrator finds appropriate execution resources for the individual tasks in a workflow, and handles task execution and monitoring. Different brokering strategies are implemented to find the most suitable resource for each workflow step.

Data Management

The workflow system can interface to data management solutions. Out of the box, a location mapping service is used to deal with file locations in an abstract manner and provide a „name space“ for data files within a workflow.

Extensibility

The Workflow Engine allows to use domain-specific workflow languages by plugging a custom converter into the Workflow Engine. The processing logic can be extended with custom features and functionality. Customized brokering strategies can be plugged into the Service Orchestrator.

Further Reading

www.unicore.eu/documentation/architecture

UNICORE

www.unicore.eu